



UNIVERSITEIT VAN AMSTERDAM

Research Project 1

Peeling Google Public DNS Onion

ANALYZING CACHE COHERENCY AND LOCALITY OF GOOGLE PUBLIC
DNS

Tarcan Turgut & Rohprimardho

Under supervision of Roland M. van Rijswijk-Deij from SURFnet

4 February 2015

Research Questions

Is there a single shared cache?

- Does the authoritative name server receive more than one query?
- Is there any delay while distributing the cache entry to other locations?
- Is level 1 cache identical?
- Does Google Public DNS respect the TTL set by the authoritative nameserver?

Where is the query to the authoritative name server coming from?

Google Public DNS

DNS

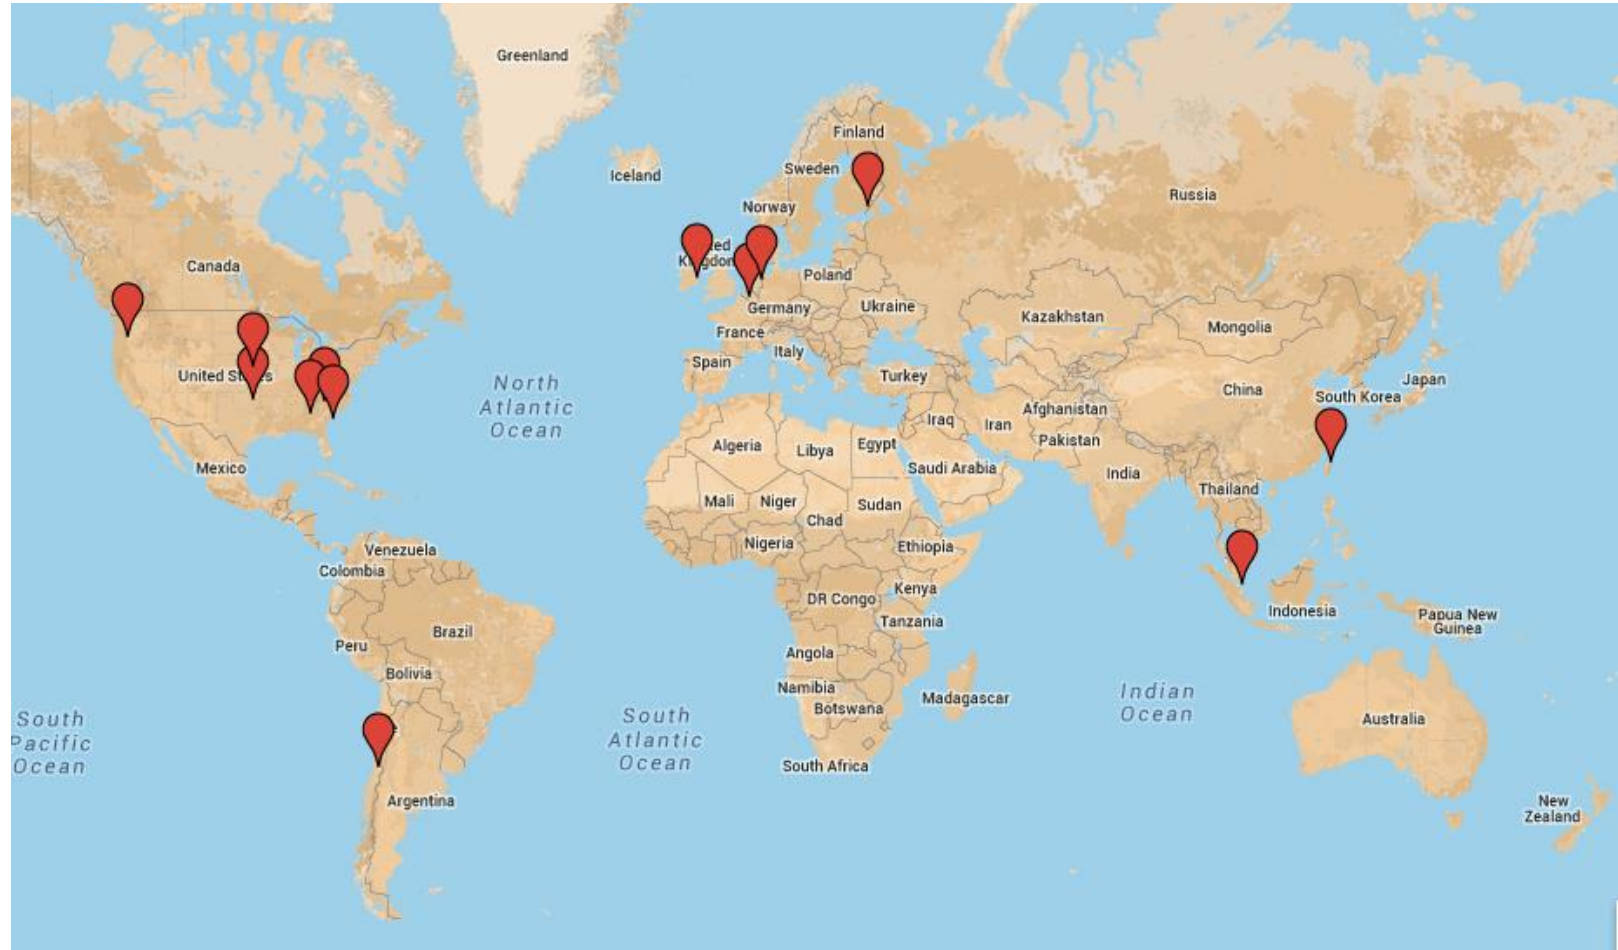
- Alternative for DNS provider

Location

- Anycast routing
- AS15169

Cache

- 2 levels
- Popular and unpopular domain names



General Topology

BIND

RIPE Atlas probes



General Topology



Source: RIPE Atlas website

Origin of the DNS Queries

Mapping the flow of the query

Use RIPE Atlas probes to send DNS queries

- Check the source of the query in the log
- 1 probe each country

Conclusion

- Query originates in Google Public DNS server close to the client
- Hints: no global single shared cache around the world

Probe Location	Query Source
Bangladesh	Singapore
Saudi Arabia	Belgium
Argentine	Chile
Ecuador	USA
Canada	USA
Algeria	Belgium
South Africa	Belgium
Finland	Finland
The Netherlands	Belgium
Russia	Finland

Round Trip Time

Compare RTT between two areas to see possible performance penalty

Traceroute to 8.8.8.8

- Southeast Asia and Western Europe (each 5 countries)
- 5 randomly picked RIPE Atlas probes

Latency is an order of magnitude higher in Southeast Asia than in Western Europe

Country Name	Average RTT (in ms)
Indonesia	17
Phillipines	45
Vietnam	40
Singapore	3
Malaysia	64
The Netherlands	5
France	3
Germany	2
Switzerland	2
Luxembourg	25

Edge Router to AS15169

To see if they all use the same edge router and if the query also came from the same origin

Same setup as the previous

- Southeast Asia and Western Europe (each 5 countries)
- 5 randomly picked RIPE Atlas probes
- Traceroute to 8.8.8.8 and also send DNS query

Result

- Edge router differs based on which RIPE Atlas probes were used
- The query not always came from the same location

Edge Router to AS15169

Conclusion

- Anycast
- Google has some kind of mechanism that takes care of the query inside AS15169

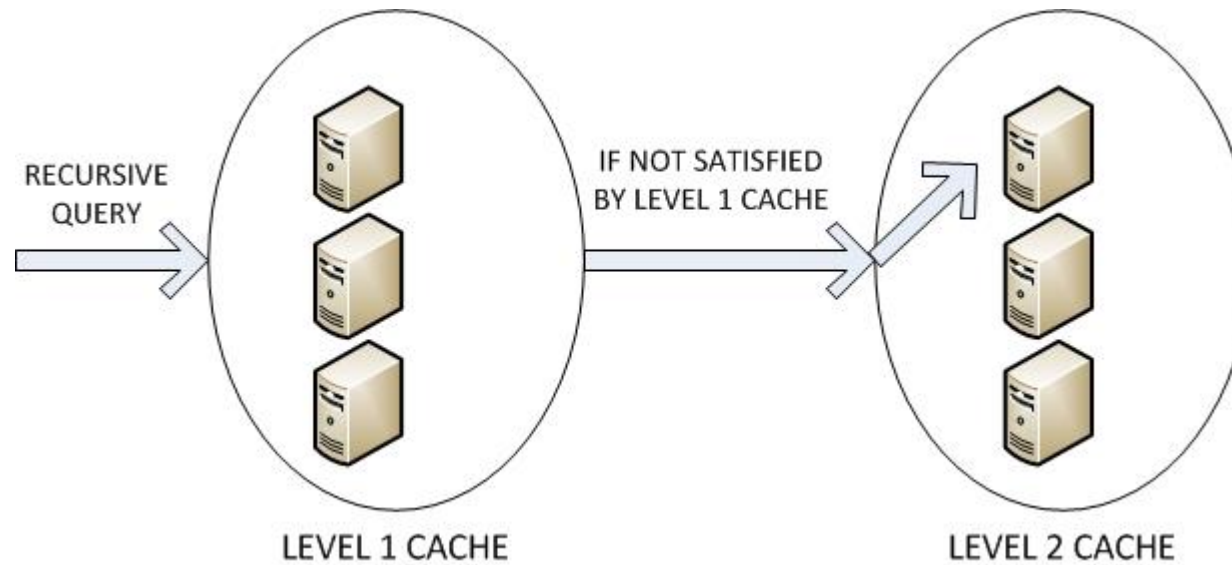


Two Levels of Caching

Level 1 cache – Most popular names (a small per-machine cache)

Level 2 cache – Unpopular names (partitioned by names)

Each level contains a pool of machines



Is Level 1 cache identical per location?

Flush Cache Tool

Flush cache for a domain

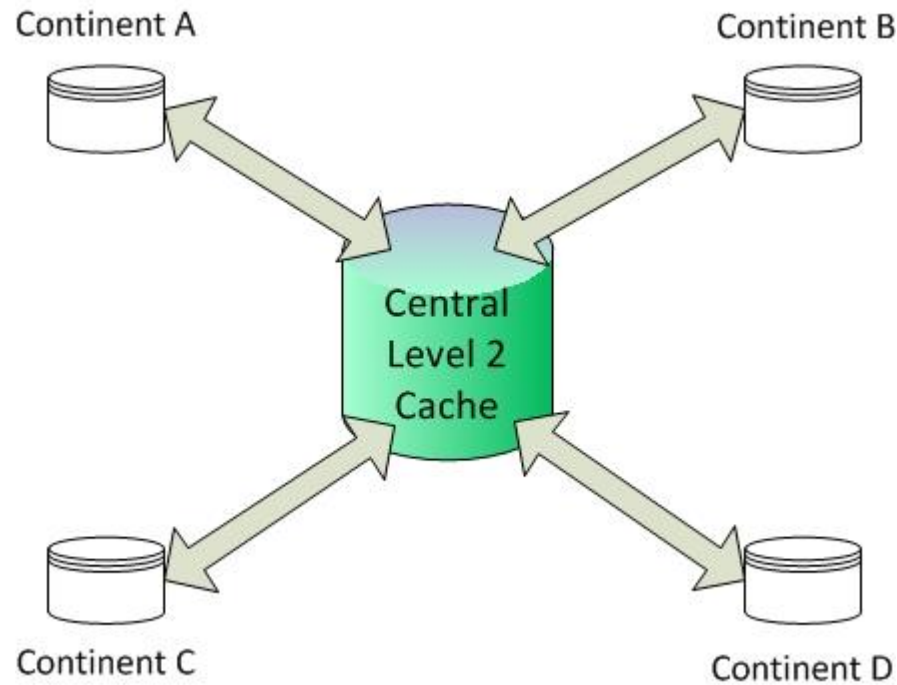
[Privacy & Terms](#)

Domain Name:

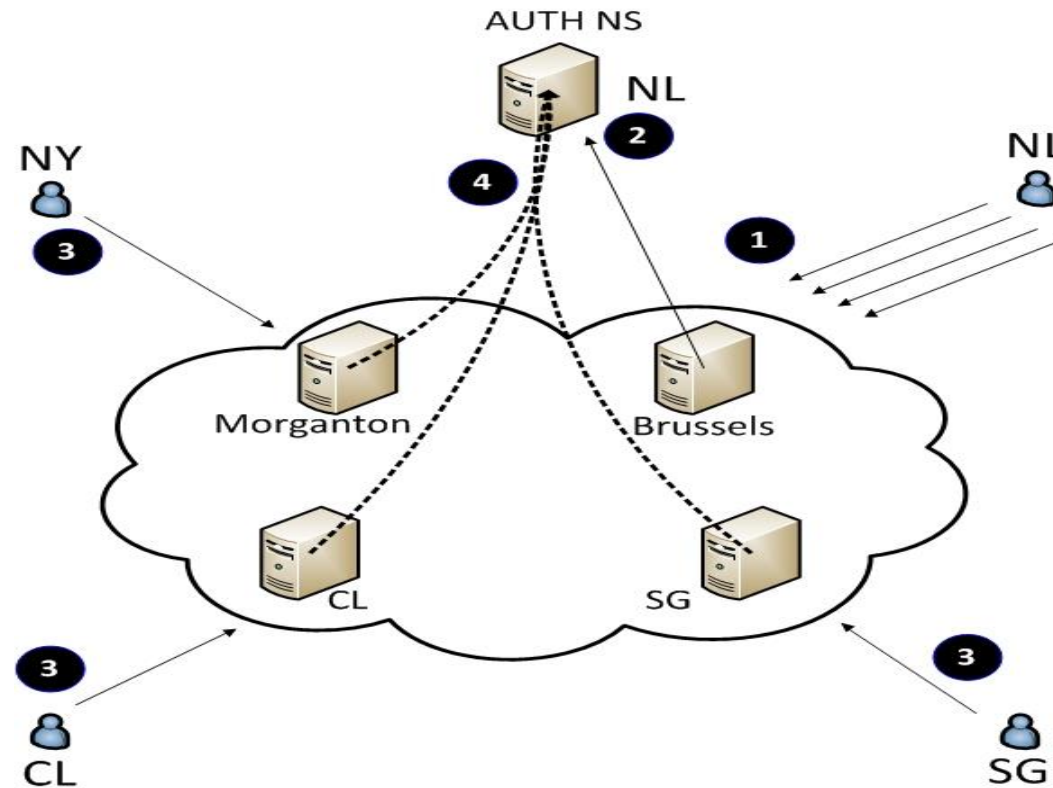
RR Type:

- Bug! They are working on it!

Global Coherency of Level 2 Cache



Global Coherency of Level 2 Cache



Result: There is NOT a single globally shared cache.

Does Google respect TTL set by authoritative name servers?

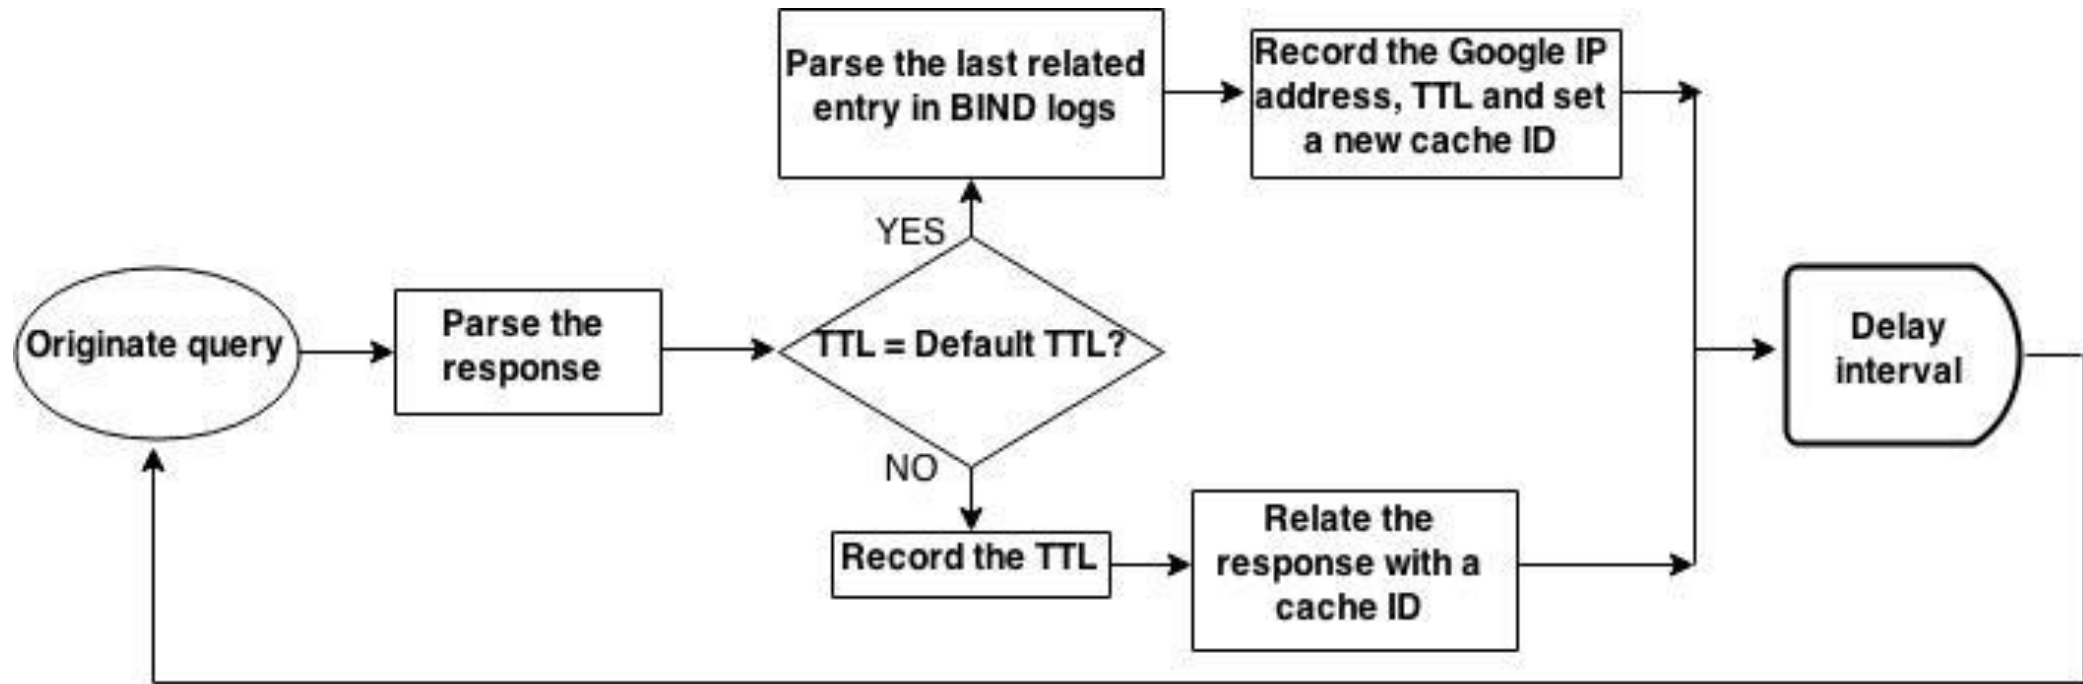
Google does NOT modify TTL values unless it is more than 6 hours

An answer for an A record with default TTL set to 1 day (86400 secs):

;; ANSWER SECTION:

day.uk.inspectorgoogle.net. 21599 IN A 178.62.38.140

Level 2 cache coherency in a single location



Level 2 cache coherency in a single location

Query ID	Timestamp	Cache ID	Google Resolver IP	TTL
1	01:50:02	1	2a00:1450:400c:c05::153	300
2	01:50:12	2	74.125.181.83	300
3	01:50:22	1	Cache Response	280
4	01:50:32	2	Cache Response	280
5	01:50:42	2	Cache Response	270
6	01:50:52	3	2a00:1450:400c:c05::153	300
7	01:51:02	2	Cache Response	250
8	01:51:12	2	Cache Response	240
9	01:51:22	1	Cache Response	220
10	01:51:32	3	Cache Response	260
11	01:51:42	4	74.125.17.209	300

Level 2 cache coherency in a single location

Finding: TTL values decrease gradually till very low values

Query ID	Timestamp	Cache ID	Google Resolver IP	TTL
1	01:50:02	1	2a00:1450:400c:c05::153	300
9	01:51:22	1	Cache Response	220
21	01:53:22	1	Cache Response	100
26	01:54:13	1	Cache Response	50
30	01:54:53	1	Cache Response	10

Implication: Google does not evict RRs from cache before TTL expires

- Cache is big enough

Level 2 cache coherency in a single location

Finding: There seems more than 1 cache in a single location.

Query ID	Timestamp	Cache ID	Google Resolver IP	TTL
1	01:50:02	1	2a00:1450:400c:c05::153	300
2	01:50:12	2	74.125.181.83	300
3	01:50:22	1	Cache Response	280
4	01:50:32	2	Cache Response	280

Implication: Level 2 cache is fragmented as opposed to Google's statement.

Level 2 cache coherency in a single location

Finding: The cache responses are coming from multiple caches.

Cache ID	Occurence
1	10
2	11
3	3
4	2

Implication: Possibly behind a load-balancer

- Not found a regular pattern pointing an algorithm such as round-robin

Level 2 cache coherency in a single location

Finding: 1st and 6th queries are handled by the same Google resolver IP

Query ID	Timestamp	Cache ID	Google Resolver IP	TTL
1	01:50:02	1	2a00:1450:400c:c05::153	300
6	01:50:52	3	2a00:1450:400c:c05::153	300

Implication: “Egress IP addresses are shared by multiple resolver” [says Google]

- A mapping between resolver IP and the cache is N/A

Level 2 cache coherency in a single location

Finding: Ghost cache

Query ID	Timestamp	Cache ID	Google Resolver IP	TTL
1	07:20:01	1	74.125.181.86	300
2	07:20:11	1	Cache Response	290 ▲ (300)
3	07:20:21	2	74.125.181.80	300
4	07:20:31	3	74.125.47.83	300
5	07:20:41	4	74.125.47.80	300
7	07:21:01	Unknown	Cache Response	250
24	07:23:52	Unknown	Cache Response	80

Implication: Not available. Extra information needed by Google!

Conclusion

The queries to an authoritative name server originates in the Google datacenter where the query is received

Not a globally centralized Level 2 cache. Expensive!

Fragmented Level 2 cache in a single location may increase the cache miss rate, consequently the response time

Level 2 cache behavior seems the same and our results are similar in different locations of Google, TTL values, frequency of originating query and time-of-day

Future Work

Hints of possible performance penalty. (Google vs. Local resolvers)

Need more information to deduce further

- Google: “We cannot disclose technical details”

Questions?